NASA/TM–20220014544

A Rideshare Tensegrity Rover Concept to Explore Titan's Lands and Oceans

Nicholas Deitrich
CRASH Lab, University at Buffalo – The State University of New York, Buffalo, New York

Karen Mae Baldonado
City College of New York, New York City, New York

Ahsan Khan
CRASH Lab, University at Buffalo – The State University of New York, Buffalo, New York

Joshua Cook
Oregon State University, Corvallis, Oregon

Louis Rizzo
CRASH Lab, University at Buffalo – The State University of New York, Buffalo, New York

Kevin Schroeder
Virginia Polytechnic Institute and State University, Blacksburg, Virginia

Javid Bayandor
CRASH Lab, University at Buffalo – The State University of New York, Buffalo, New York

Jamshid Samareh
Vehicle Analysis Branch, NASA Langley Research Center, Hampton, Virginia

October 2022

NASA STI Program Report Series

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA scientific and technical information (STI) program plays a key part in helping NASA maintain this important role.

The NASA STI program operates under the auspices of the Agency Chief Information Officer. It collects, organizes, provides for archiving, and disseminates NASA's STI. The NASA STI program provides access to the NTRS Registered and its public interface, the NASA Technical Reports Server, thus providing one of the largest collections of aeronautical and space science STI in the world. Results are published in both non-NASA channels and by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA Programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counterpart of peer-reviewed formal professional papers but has less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services also include organizing and publishing research results, distributing specialized research announcements and feeds, providing information desk and personal search support, and enabling data exchange services.

For more information about the NASA STI program, see the following:

- Access the NASA STI program home page at http://www.sti.nasa.gov

- Help desk contact information:

https://www.sti.nasa.gov/sti-contact-form/
and select the "General" help request type.

NASA/TM–20220014544

A Rideshare Tensegrity Rover Concept to Explore Titan's Lands and Oceans

Nicholas Deitrich
CRASH Lab, University at Buffalo, Buffalo, NY

Karen Mae Baldonado
City College of New York, New York, NY

Ahsan Khan
CRASH Lab, University at Buffalo, Buffalo, NY

Joshua Cook
Oregon State University, Corvallis, OR

Louis Rizzo
CRASH Lab, University at Buffalo, Buffalo, NY

Kevin Schroeder
Virginia Polytechnic Institute and State University, Blacksburg, VA

Javid Bayandor
CRASH Lab, University at Buffalo, Buffalo, NY

Jamshid Samareh
Vehicle Analysis Branch, NASA Langley Research Center, Hampton, VA

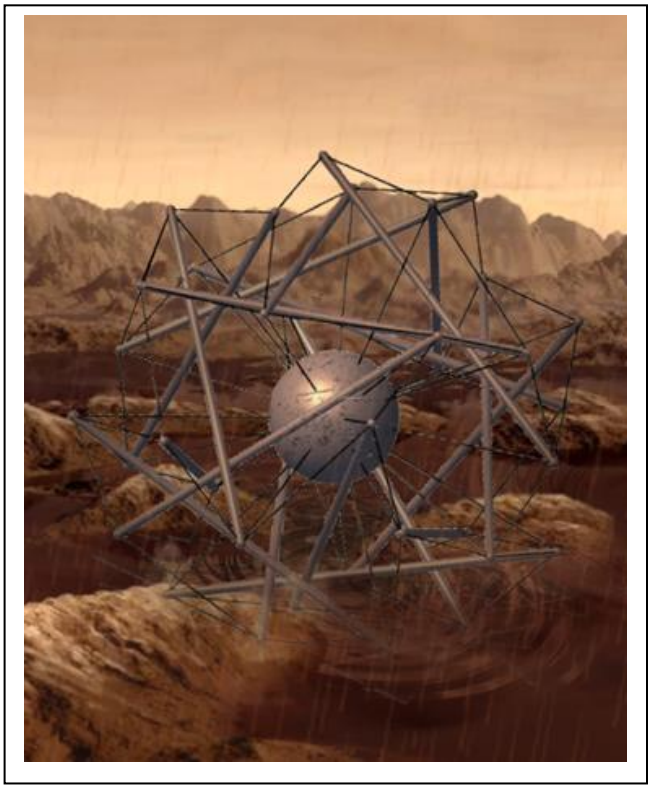

National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681-2199

October 2022

Available from:

NASA STI Program / Mail Stop 148
NASA Langley Research Center
Hampton, VA 23681-2199
Fax: 757-864-6500

Acknowledgments

The Phase II funding under the NASA Innovative Advanced Concepts (NIAC) Program is acknowledged and appreciated. We would like to thank the *CRASH* Lab of the University at Buffalo - The State University of New York, as the principal investigator organization for the TANDEM program; the NASA Langley Research Center; NASA Ames Research Center, for their partnership throughout the TANDEM development; Massimo Vespignani and Jonathan Bruce from NASA Ames for their involvement in the project, invaluable contributions, and strong support of the team; Jessica Gangitano of NASA Langley University Space Research Association for her expert and patience guidance through the internship process; Vehicle Analysis Branch at Langley Research Center for hosting the first four authors; and Robotics Branch at Ames Research Center for hosting the fifth author.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

NOMENCLATURE

A	area (m^2)
ADEPT	Adaptable Deployable Entry and Placement Technology
b	number of rigid elements
C_b	bar element connectivity matrix
C_D	drag coefficient
C_s	cable connectivity matrix
CNC	Computer Numerical Control
d(g)	distance to goal (m)
DoF	degrees of freedom
EDL	Entry, Descent, and Landing
EDL-L	Entry, Descent, Landing, and Locomotion
ESA	European Space Agency
F	fitness
HL	total heat load (J/m^2)
ICoSS	Integrated Composite Stiffener Structure
l_b	length of rigid element (m)
$\boldsymbol{l_s}$	input cable lengths vector (m)
m	mass (kg)
MER	Mars Exploration Rover
MSR	Mars Sample Return
NASA	National Aeronautics and Space Administration
N_{goals}	number of goal positions
N_{nodes}	number of nodes
N_{pairs}	number of cable actuation pairs
PICA	Phenolic Impregnated Carbon Ablator
R	reward function
s	number of cables
SLA	Super Lightweight Ablator
TANDEM	Tension Adjustable Network for Deploying Entry Membrane
TH	TPS thickness (m)
TPS	Thermal Protection System
TRL	Technology Readiness Level
V	velocity (m/s)
x	nodal x positions (m)
y	nodal y positions (m)
z	nodal z positions (m)
$\mathbf{x_{FF}}$	nodal x positions from form finding (m)
$\mathbf{y_{FF}}$	nodal y positions from form finding (m)
$\mathbf{z_{FF}}$	nodal z positions from form finding (m)
x_i,	i^{th} node x position (m)
y_i	i^{th} node y position (m)
z_i	i^{th} node z position (m)
β	ballistic coefficient (kg/m^2)

SUMMARY

The Tension Adjustable Network for Deploying Entry Membrane (TANDEM) concept is a phase II NASA Innovative Advanced Concepts (NIAC) project. One of the goals for this project was to extend the application of tensegrity-based rovers to a rideshare concept for a Titan mission. This report provides preliminary details of the TANDEM concept for Titan. System analyses of the entry vehicle were performed and compared to data from the Huygens mission to Titan. The entry aeroshell was based on Huygens and Stardust aeroshells. Modeling, simulation and design of the tensegrity deployment were performed using a nonlinear optimization form-finding algorithm. Lastly, a neuroevolution-based machine learning control strategy was applied, which produced efficient rolling locomotion gaits.

1. Introduction

1.1. Background

The *CRASH* Lab and NASA Langley developed TANDEM, a planetary rover concept in which all aspects of entry, descent, landing, and locomotion (EDL-L) are achieved by a single, multifunctional tensegrity structure [1-2]. The tensegrity rover is comprised of rigid elements (e.g., bars) connected to a network of tension elements (e.g., cables) and attached to a central payload module. By actuating the cables, TANDEM morphs its shape to be stowed inside an aeroshell and deployed into its landing configuration. For locomotion, it shifts its center of gravity to roll across terrains. Loads are distributed throughout the tension network to protect the payload at impact. The lightweight design, omnidirectional landing, and adaptive locomotion capabilities make TANDEM a strong candidate to explore extreme environments on most celestial bodies in our Solar System. Venus and Titan are of particular interest to the TANDEM Team, since both planets have thick atmospheres, limiting the exploration of the surface topography from orbit. TANDEM was originally proposed for a Venus surface exploration mission. This report details the application of TANDEM as a rideshare concept candidate to explore Titan's lands and oceans.

1.2. Concept of Operations

The original TANDEM concept combined the tensegrity rover with an Adaptable Deployment and Entry Placement Technology (ADEPT) [3] semi-rigid heatshield to create a novel multifunctional EDL-L vehicle. This deployable aeroshell concept is suitable for Titan's atmosphere, but this study selected the flight-proven, rigid aeroshells that provide a higher technology readiness level (TRL) and fit the timeline for upcoming Titan rideshare opportunities. The entry, descent, and landing (EDL) concept is similar to the Huygens mission. Two flight proven heatshields were evaluated: Huygens and Stardust. A modified Huygens aeroshell with Mars Exploration Rover (MER) backshell was also evaluated.

Figure 1 shows the concept of operations of TANDEM exploring Titan's lands and oceans. After peak heating in the upper atmosphere, a parachute pulls out the aeroshell, followed by TANDEM's free-fall to the surface. The selected design parameters keep the impact velocity within the designed capabilities of TANDEM, eliminating the need for a second parachute. The spherical landing configuration protects the payload module at any impact orientation. This is critical because of Titan's unknown topography. Following impact, TANDEM will explore the surface by actuating its cables, causing a rolling locomotion. In addition to rolling locomotion, TANDEM could use its actuated tensegrity

frame to swim across the liquid bodies on Titan. One concept is to have elliptical cross sections on the rigid elements and use the flatter end as a paddle to swim with a breaststroke-like gait.

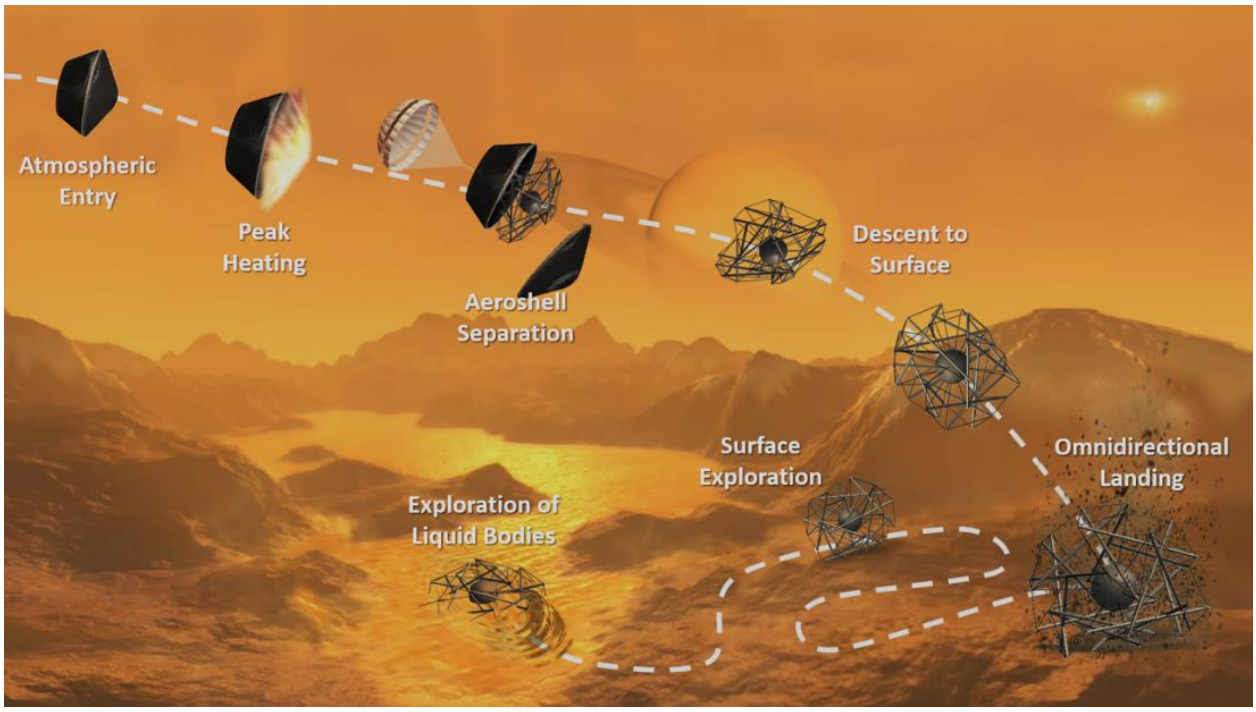

Figure 1: Concept of Operations of TANDEM on Titan.

2. TANDEM Trajectory

TANDEM's baseline design for a mission to Titan had a similar ballistic coefficient to Huygens. The systems analysis followed the approach of a previous study [4]. The EDL segment of the mission was modeled using 3 degrees-of-freedom (DoF) equations of motion solved with a Python module called Quick EDL. Quick EDL provided altitude, velocity, acceleration, and stagnation point heating.

3. Aeroshell Design

TANDEM's original concept mission to Venus uses a deployable heatshield, which is actuated by the rover's tensegrity frame. This proposed concept mission to Titan evaluated two rigid and flight-proven aeroshells with higher technology readiness levels (TRL) than the deployable aeroshell of the original TANDEM. The higher TRL is an important consideration for a potential rideshare with the Dragonfly Titan mission scheduled to launch in 2027.

3.1. Huygens and Stardust Aeroshells

Two flight proven aeroshell concepts were considered: Huygens and Stardust. Figure 2 shows TANDEM packed within each aeroshell. The Huygens aeroshell successfully entered the Titan atmosphere at 6 km/s and landed. Stardust was a return capsule sent back to Earth after gathering comet samples and cosmic dust. It entered Earth's atmosphere at 12.8 km/s and landed in the U.S. Air Force Utah Test and Training Range. Both Earth and Titan have nitrogen rich atmospheres, which makes Stardust's 60° sphere cone aeroshell a viable option for a Titan mission. The Stardust backshell also has a larger volume that can hold a scaled up TANDEM. Each aeroshell is designed differently, most notably the Huygens backshell is narrower than the heatshield, to lower ballistic coefficient and improve the dynamic stability during the Titan atmospheric entry. The TANDEM payload has a much lower center of gravity, mitigating some of the entry dynamic stability issues.

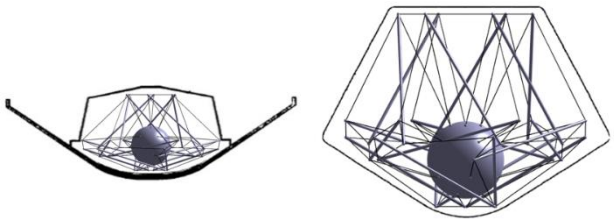

Figure 2: TANDEM Inside Huygens Aeroshell (Left) and Stardust Aeroshell (Right).

The TANDEM Titan aeroshell design uses the shape of the Stardust aeroshell with the ballistic coefficient of the Huygens aeroshell. Ballistic coefficient is defined as:

$$\beta = \frac{m}{C_D A} \tag{1}$$

Using Huygens' ballistic coefficient of 38 kg/m^2 and flight path angle of 65° results in the same trajectory and heating profile. To maintain the same ballistic coefficient, the ratio of mass to area was kept constant.

3.2. Thermal Protection System (TPS)

The estimated heat load and entry velocity were used to calculate the TPS thickness using surrogate models developed by Sepka and Samareh [5]. Based on maximum heat flux, Phenolic Impregnated Carbon Ablator (PICA) was selected for the heat shield and Super Lightweight Ablator (SLA) for the backshell. PICA and SLA are proven technologies and were used for the Stardust aeroshell [6].

Equations (2) and (3) were used to estimate the TPS thickness for PICA and SLA, respectively.

$$TH = 1.8696\left(\frac{HL}{V^2}\right)^{0.1873} \tag{2}$$

$$TH = 0.0064\left(\frac{HL}{V^2}\right)^2 + 0.0961\left(\frac{HL}{V^2}\right) + 0.3322 \tag{3}$$

A uniform TPS thickness was used to calculate the total TPS mass.

3.3. Aeroshell Structure

The Integrated Composite Stiffener Structure (ICoSS) concept was proposed for the future Mars Sample Return (MSR) mission [7]. It is comprised of a lightweight foam core that forms the stiffener geometry going across the structure, shown in red in Figure 3. The stiffeners are connected by co-cured junctions that allow for efficient load distribution throughout the structure during launch and entry. The Computer Numerical Control (CNC) machined foam core offers the benefit of design flexibility to be adapted to lander missions.

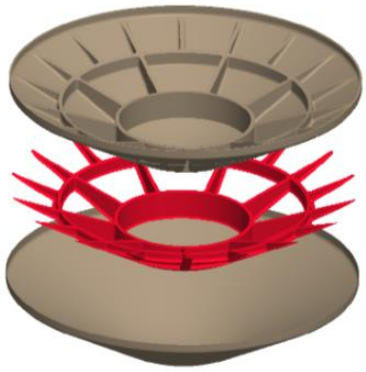

Figure 3: (Top) Heatshield, (Middle) Stiffener Structure, (Bottom) Outer Skin [7].

The ICoSS concept will be redesigned to accommodate future design requirements for TANDEM. The concept will be optimized to sustain loads from launch and entry. Current work is focused on performing size optimization of the ICoSS forward shell.

The original ICoSS concept was designed to withstand surface impact for MSR. However, TANDEM will not require impact energy absorption by its forward shell. Future work will investigate redesigning the system to withstand only aerodynamic loading at a reduced mass.

3.4. Master Equipment List

The TANDEM lander masses reported by Schroeder, Samareh and Bayandor [1] were scaled for the Titan mission. Table 1 shows the master equipment list for TANDEM and its aeroshell scaled to the diameter of the Stardust aeroshell, with the ballistic coefficient of the Huygens aeroshell.

Table 1: Master Equipment List for TANDEM Titan Mission

	TANDEM		
	Current Best Estimate Mass (kg)	Growth	Predicted Mass (kg)
Total Mass	269		344
Aeroshell	102	0.3	131
Structures	33		43
Heatshield	11	0.3	14
Structural Backshell	22	0.3	29
TPS	69		89
Heatshield	24	0.3	32
Backshell	45	0.3	58
Lander	167		213
Scientific Payload	10	0.3	12
Thermal	17	0.3	22
Power	3	0.3	4
Comms. Avionics & Electronics	6	0.3	8
Secondary Structure	30	0.3	38
Landing System	100	0.3	129

4. Tensegrity Modeling

Tensegrity-based rovers are lightweight, redundant, and packable platforms particularly suitable for planetary exploration missions. The SUPERball tensegrity rover demonstrated packability, payload protection from impacts, and robust surface locomotion at the NASA Ames Roverscape test facility [8]. The TANDEM concept adds a higher degree of versatility, controllability, and redundancy to the system [1]. SUPERball is comprised of 6 rigid elements and 24 outer cables. The baseline TANDEM concept has 18 rigid elements, with 54 outer cables and 24 inner cables supporting the payload module. There are a total of 78 cables. The added controllability allows TANDEM to precisely morph to various aeroshell geometries and perform more rapid locomotion with less energy. The primary tool used to design, simulate, and control the tensegrity structure is the MATLAB Tensegrity Dynamics Modeling and Simulation Framework developed at NASA Ames [9].

4.1. Tensegrity Topology

The topology of a tensegrity structure describes how the rigid elements and cables are connected to one another. The topology is represented by a connection matrix, which contains rows corresponding to each rigid element and cable, and columns corresponding to end nodes. The matrix is sparse with each row containing a 1 and -1, signifying the nodes from which that element starts and ends. Two distinct topologies were originally proposed for the TANDEM concept [1]. The fundamental rigid elements and cables of these two topologies are labeled in Figure 4.

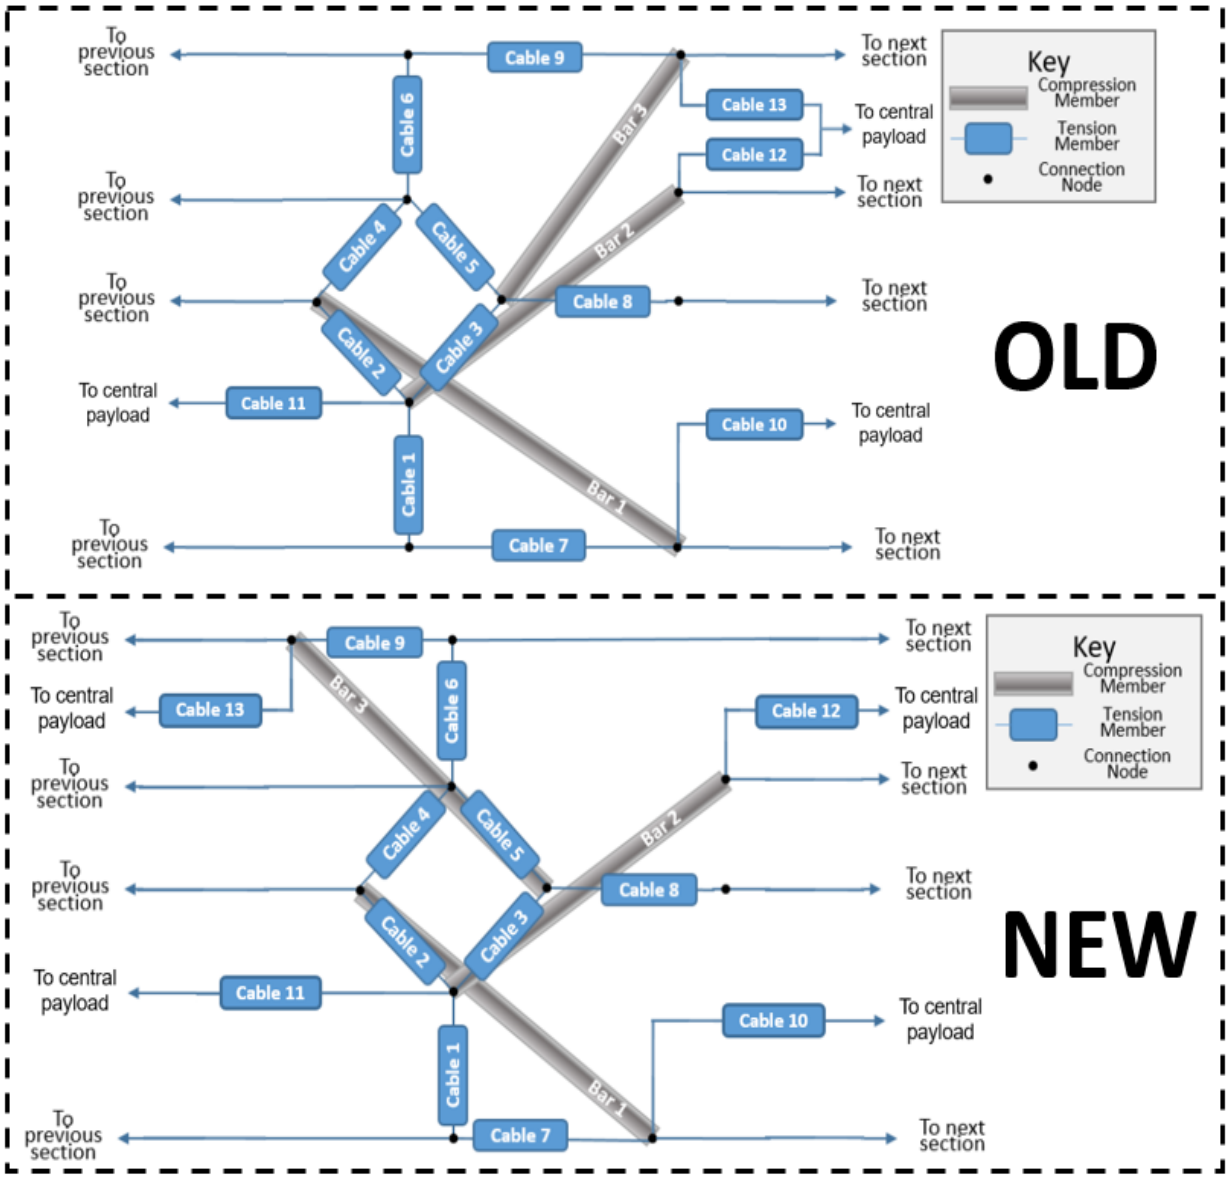

Figure 4: Original and Proposed Topologies for TANDEM (Adapted from [1]).

The orientation of the upper bars differ between the two topologies. Figure 5A shows the two tensegrity structures with their fundamental bars highlighted. The positions of the nodes of the two tensegrity structures are the same, but their connections differ. The omnidirectional landing capability of each topology was assessed by simulating all possible orientations of impact and calculating the maximum cable tension and the maximum deceleration experienced by the payload. The x and y axes of Figures 5B and 5C correspond to spherical coordinates (i.e., longitude and latitude) at which the structure impacts the ground. Figure 5 shows the results for an impact velocity of 15 m/s, payload mass of 50 kg, and 0.75 kg point masses attached to each node. These values reflect a benchmark system design and do not represent the TANDEM Titan mission design. The maximum cable tension, seen in Figure 5B, does not vary significantly between the two topologies, however the maximum payload deceleration, shown in Figure 5C, is significantly different. The dark red regions in the old topology were identified as orientations resulting in high impact loads compared to the new model. The maximum deceleration decreased from 292 Earth g's for the original topology to 279 Earth g's for the proposed topology. The orientations corresponding to maximum payload deceleration in the

original topology are exclusively found in the southern hemisphere, whereas the proposed topology is more evenly distributed.

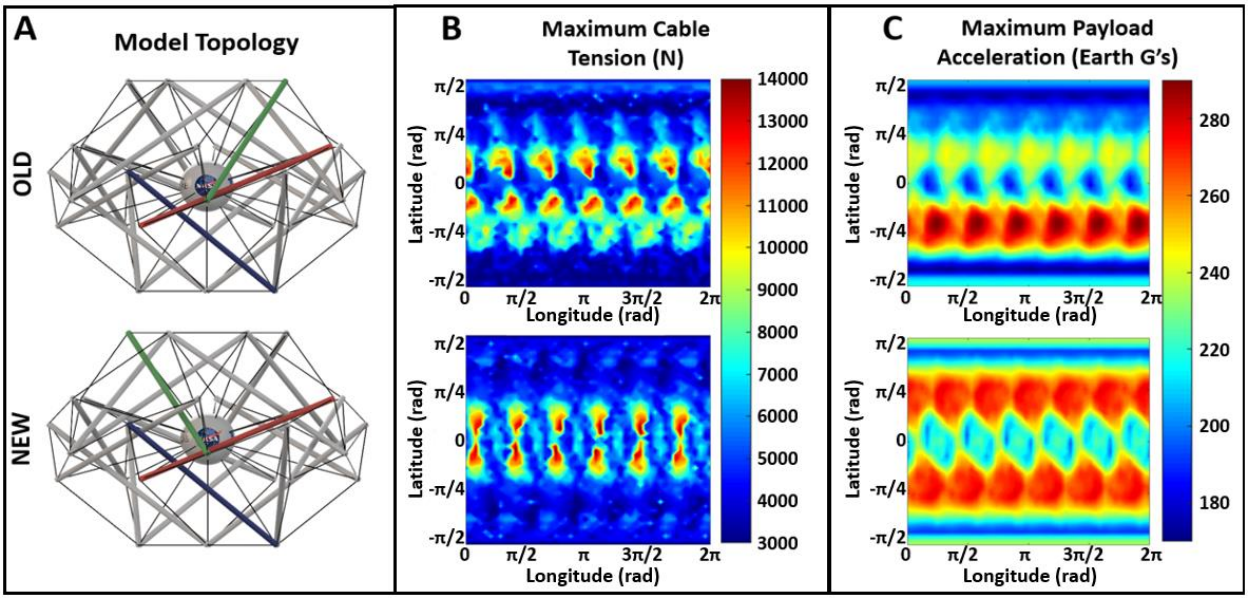

Figure 5: Comparison of Original and Proposed Topologies.

The new baseline connectivity was used to generate new tensegrity structures with different numbers of bars. Trade-offs between the structural mass, payload protection, and locomotive capabilities were investigated and compared to the baseline 18-bar model. Tensegrity rovers with a fewer number of bars may be better suited for different mission requirements, such as companion rovers for lunar astronauts. Figure 6 shows tensegrity rovers with 9, 12, and 24 bars with the new baseline topology.

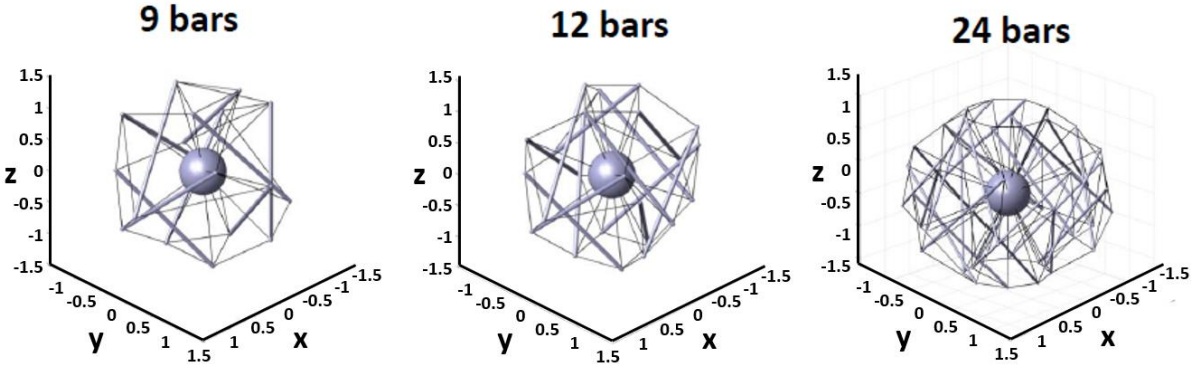

Figure 6: Tensegrity Models.

4.2. Form Finding

Form-finding algorithms are used to determine the tensegrity configurations, both while packed inside the aeroshell and when deployed. Several form-finding algorithms exist [10], of which the kinematical approach of nonlinear optimization was chosen to create new configurations of the tensegrity structure. The algorithm determines the optimal shape and loading of the tensegrity structure by maximizing the lengths of the rigid elements, subsequently minimizing

the lengths of the cables. The input parameters are the connection matrix and normalized lengths of each cable, and the algorithm determines the optimal nodal coordinates.

Equations (4) describe the form-finding algorithm, where the objective function is to maximize the length of one of the rigid elements. The first set of constraints enforce the cables to be the input length, and the second set of constraints keep the lengths of the rigid elements equal. There is one nonlinear constraint equation for each cable and rigid element in the structure, resulting in a set of 96 nonlinear constraint equations for TANDEM. The problem was solved using the constrained minimization routine "fmincon" in MATLAB*. The algorithm outputs the nodal coordinates of the tensegrity structure with the prescribed cable lengths and equal rigid element lengths. These coordinates are then scaled for the desired size of the tensegrity structure.

$$\text{Minimize:} \quad -(l_b)^2 = -\left((x_i - x_j)^2 + (y_i - y_j)^2 + (z_i - z_j)^2\right)$$

$$\text{Subject to:} \quad \begin{cases} (C_s\boldsymbol{x})^2 + (C_s\boldsymbol{y})^2 + (C_s\boldsymbol{z})^2 - (\boldsymbol{l_s})^2 = \mathbf{0} \\ (C_b\boldsymbol{x})^2 + (C_b\boldsymbol{y})^2 + (C_b\boldsymbol{z})^2 - (l_b)^2 = \mathbf{0} \end{cases} \tag{4}$$

The sum of the first three terms in each of the constraint equations represents the lengths of the cables and bars. Using the kinematic form-finding algorithm, TANDEM's shape is morphed by changing the prescribed lengths of the cables $\boldsymbol{l_s}$. Table 2 shows the input parameters for configurations generated with the form-finding algorithm. Empty cells indicate cables whose length was not prescribed during form finding. Cables 1-13 are the fundamental, axisymmetric cables in the tensegrity structure and are labeled in Figure 4. Cables 10 through 13 control the position of the central payload module. Figure 7 shows selected configurations.

Table 2: Normalized Cable Lengths for Various TANDEM Configurations [1].

Configuration	Cable Length Coefficients												
	1	2	3	4	5	6	7	8	9	10	11	12	13
Stowed	2.40	3.50	1.50	1.50	3.50	2.40	1.00	3.00	1.00				
Deployed	2.85	1.00	1.00	1.00	1.00	4.00	1.00	3.00	1.00				
Descent	4.00	1.00	1.00	1.00	1.00	2.85	1.00	3.00	1.00				
Default	1.50	1.50	1.00	1.00	1.50	1.50	1.00	1.50	1.00				
Stardust	3.00	1.00	1.00	1.00	1.00	3.00	1.00	3.00	2.00	1.50	3.00	3.50	4.00
MER Backshell	3.00	1.00	1.00	1.00	1.00	4.25	1.00	3.00	1.00	1.25	3.50	4.50	3.00

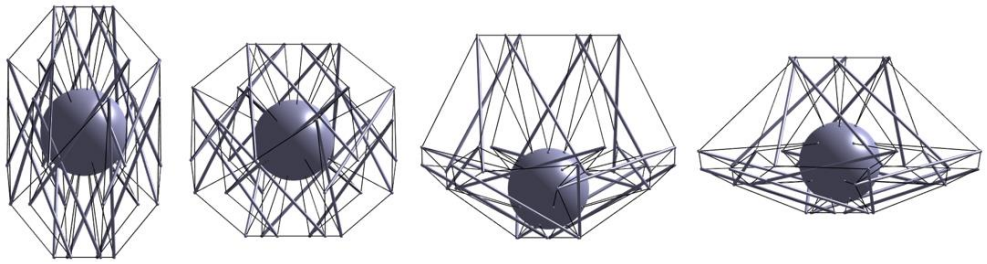

Figure 7: TANDEM Configurations. (Left to Right: Stowed, Default, Stardust, MER Backshell)

*Reference to or appearance of any specific commercial products, processes, or services by trade name, trademark, manufacturer, or otherwise does not constitute or imply its endorsement, recommendation, or favoring by the U.S. Government or NASA.

MathWorks https://www.mathworks.com/products/matlab.html

4.3. Deployment Sequence

A key feature of the TANDEM concept is the rover's ability to deploy from its packed shape inside the aeroshell to its unpacked shape using the tensegrity-actuated frame. By linearly varying the length coefficients from one configuration to another, a purely kinematic deployment sequence was created using the form-finding algorithm. Although the deployment captures the shape change and preserves the rigid elements' lengths, the mass of the bars and payload are ignored. To capture the effect of the structural mass on the deployment sequence, a controller was created in the MATLAB simulation environment. This controller was generated by a genetic algorithm, which uses the output nodal positions from the form-finding algorithm to transform TANDEM into a desired configuration.

Equation (5) is the fitness function that represents the sum of the differences between the current simulation and desired form-finding node positions. To maximize fitness, the function is multiplied by -1, so that the nodal positions of the most fit selection are closest to the desired configuration. The terms $\boldsymbol{x}$, $\boldsymbol{y}$, $\boldsymbol{z}$ are the node positions after the simulation, and $\boldsymbol{x_{FF}}, \boldsymbol{y_{FF}}, \boldsymbol{z_{FF}}$ are the desired node positions from the form-finding algorithm.

$$F = -\text{sum}((\boldsymbol{x} - \boldsymbol{x_{FF}})^2 + (\boldsymbol{y} - \boldsymbol{y_{FF}})^2 + (\boldsymbol{z} - \boldsymbol{z_{FF}})^2) \tag{5}$$

Each population of the genetic algorithm contained 13 cable rest-length parameters, one for each unique cable. The cable rest-length is the parameter which controls the cable actuation. The first generation included randomly selected rest-lengths, which were simulated and then evaluated using the fitness function in Equation (5). Populations were ranked using a binary tournament, where the winner continues to the next generation unchanged, and the loser becomes a mutated version of the winner. The mutation contained 13 elements, where each element had a 10% chance to mutate. The mutation was chosen from a random normal distribution with a mean of 1 and a standard deviation of 0.05. The cable rest-lengths were multiplied by the mutation, such that 10% of these rest lengths mutated. This genetic algorithm was run with 32 populations and 1000 generations to create robust configuration changing rest-length controllers. Figure 8A shows the maximum fitness increasing across the generations until it plateaus at zero. Figure 8B shows the prescribed configuration from the form-finding method, compared to the result of the genetic algorithm in Figure 8C. There is negligible difference between these two configurations because the fitness defined in Equation (5) approached zero.

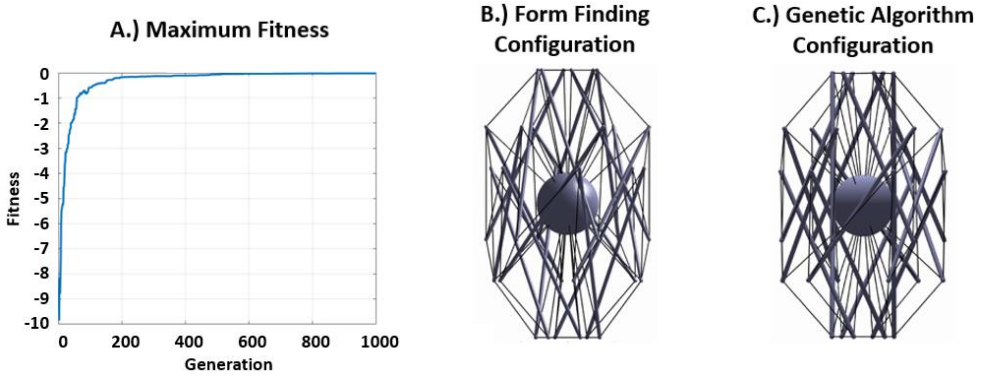

Figure 8: Genetic Algorithm Results.

The genetic algorithm accounted for the structure and payload masses during simulated shape changes. Figure 9 shows the data from the simulator, including the cable tensions, lengths, and rest-lengths. The data corresponds to the ball-to-stowed shape change described in previous paragraph with a benchmark payload mass of 50 kg, point masses of 0.75 kg at each node, 2.0 m rigid element lengths, and Earth's gravity. Figure 9A shows the simulated length of each of the 13 unique cables with a dashed line, compared to the kinematic cable lengths which are plotted as solid lines of the corresponding color. Figures 9B and 9C show the cable rest-lengths over time and the cable tension over time, respectively. The stiffness of each cable is 4000 N/m, which is the design stiffness for the cables on SUPERball [11].

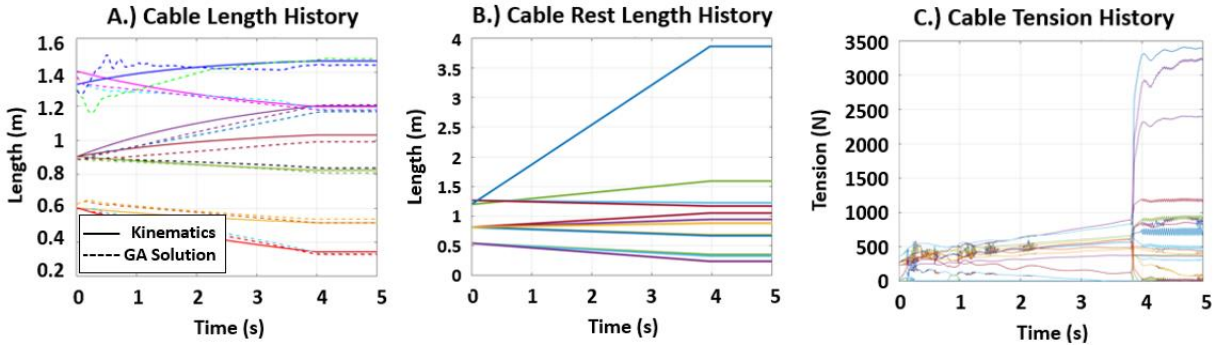

Figure 9: Ball-to-Stowed Simulation.

5. Control Strategy

A learning-based nonlinear controller was used to generate efficient locomotion gaits for TANDEM. Due to the large degree of freedom and nonlinearities inherent in tensegrity structures, creating a closed-form representation of the system dynamics is not feasible. Instead, a high-fidelity simulation was used to represent the dynamics of the system. Using this simulation, controllers learn efficient control policies. The learning-based control method selected for this application was neuroevolution. Neuroevolution generates a quality controller by evolving a neural network controller which maps a state to an action to maximize a performance score.

5.1. Neuroevolution

Neuroevolution is a method of training neural networks to map a desired input-output relationship. Due to the gradient-free nature of evolution, these networks can be trained using performance functions that are not differentiable with respect to the network parameters, such as a reward function for a controller. This makes neuroevolution a suitable method for training controllers [12]. Here, a network follows a Markov decision process, where the network must map a state to an action that maximizes the reward received [13]. Neural networks are chosen to learn this function due to their ability to learn any arbitrary function given a large enough network size [14]. Assuming the network has an achievable parameter set, a network can model a high-quality control policy [15].

These network parameters are optimized via evolution. In this study, a population of networks was used to keep a variety of solutions. Each network was evaluated using a fitness function which defined the performance of the policy in the simulation using a scalar value. Networks with higher fitness scores survive into the next generation, while networks with lower fitness values are discarded. These discarded networks are then replaced by mutated copies of the higher fitness networks. The purpose of mutation is to generate a slightly different neural network from a previous version. This is done by randomly modifying the weights and the topology of the neural network [16]. These mutated copies will potentially score better than their predecessors. If they do, they will continue to the next generation, and if not, then they will be eliminated. This concludes a generation. After a long series of small positive mutations, the population of networks is likely to contain at least one viable neural network controller [17].

5.2. Related Tensegrity Locomotion Work

Due to the complexity of tensegrity dynamics, learning-based methods of control are preferred. By evolving a discrete or sinusoidal signal for each rest-length controller, a tensegrity unit can follow a high-quality rolling motion [18], [19]. However, this type of open loop control limits the applicability of the tensegrity unit. The controllers are unable to adapt to disturbances in their environment or move to a goal position. For realistic application of these tensegrity structures, a closed loop controller is needed.

One method of providing a closed-loop control on a tensegrity unit is the Flop and Roll algorithm. This algorithm trains a tensegrity rover to perform a series of flop motions to generate a rolling pattern. Extending this algorithm to the TANDEM structure is not feasible as Flop and Roll leverages the high degree of symmetry in the SUPERball structure that is not present in TANDEM [20]. Instead, a centralized method of learning control for the TANDEM tensegrity unit was used.

5.3. Problem Formulation

In this work, the state space, action space, and reward function of the Markov decision process were defined. The state of the TANDEM rover consists of the three Euler angles of the payload orientation and a vector describing the distance toward the goal in the x and y directions. These variables were chosen as they give the rover its relative position towards its goal location. This information is sufficient for locomotion on flat ground, as used in the simulations. More information will be required for rough terrain, where the rover must incorporate additional information into the state representation, such as contact points with the ground and slope of the terrain.

The TANDEM rover was controlled via a series of rest-length controllers, one for each cable. The cable rest-length is the parameter which controls cable actuations. The action space consists of setting the desired rest-length of each of the 78 cables. The action space is large and difficult to explore since most configurations introduce slack into the cables. Learning a series of lengths that lead to a stable configuration is also difficult as there is no clear metric for the stability of a given configuration. To resolve this issue, the action space was limited so that the controller did not directly control the length of each cable. Instead, the cable lengths started at a stable configuration, i.e., the ball configuration. Then the controller selected pairs of cables, where one cable was extended by a factor of 1.5 and the other cable was contracted by a factor of 0.5. The remaining cables stayed at their original configured length. By discretizing the controller to deviate from the ball configuration, the stability of the ball configuration was preserved.

Two different reward functions were used to train two distinct behaviors. The first reward function prioritized the speed of the rover as shown in Equation (6). Here, the rover was rewarded based on the final distance to the goal. The goal was set to a distance that the rover could not reach in the allotted time. As a result, the rover learned to move quickly in the direction of the goal to maximize the reward it received.

$$R_1 = -d(g)_f \tag{6}$$

Where R_1 is reward score and $d(g)_f$ is the distance to the goal location at the final timestep. The second reward function prioritized robust control as shown in Equation (7). Here, the rover ran N_{goals} number of simulations, each with a different goal. For each simulation, the distance of the rover to goal was recorded at the beginning and the end of the simulation. The lower of these two numbers was summed to the running sum. This sum was multiplied by -1 to maximize the negative distance toward the goal.

$$R_2 = -\sum_{g=1}^{N_{goals}} \min\left(d(g)_0, d(g)_f\right) \tag{7}$$

Where R_2 is reward score, N_{goals} is the number of goal locations, $d(g)_0$ is the distance to the goal location at the initial timestep, and $d(g)_f$ is the distance to the goal location at the final timestep. The reward function in Equation (7) was chosen to provide an easier signal to learn from, when compared to only adding the final distances like Equation (6). Consider the case where the rover must learn to move to a point 1 m to the left and another point 1 m to the right. If the rover never moves, then it will receive a reward of -2 as it is 1 m from both goal locations. Later, the rover moves towards the goal on the right and sits on top of it. Using R_1, the reward is -2 as the rover is 2 m away from the left goal and 0 m away from the right goal. The rover is exhibiting a desirable behavior but is not rewarded for it. Using R_2, the rover would receive a reward of -1 as it is not penalized for moving away from the left goal. This stepping-stone reward structure is easier for the rover to learn from as it does not include any unnecessary penalties.

5.4. Simulation Setup

In each simulation, the neural networks used the same structure: 6 inputs, 42 hidden layers, and 78 output layers. The activation function for each layer was a sigmoid. The topology of the networks remained static. The weights and biases were initialized via normal distribution with a mean of 0.0 and a standard deviation of 0.05.

A population size of 32 networks was used in each simulation. The networks were selected using binary tournament selection [21]. Fitnesses were assigned using either R_1 or R_2. To mutate a network, each network parameter was mutated separately. If a parameter was chosen, a number was drawn from a standard normal distribution and summed to the parameter's previous value.

To select pairs of cables to extend and contract for each time step, the state was fed into the neural network. The output vector determined which cables to extend and contract. Each value in the output corresponds to a cable. The N_{pairs} highest values in the output have their corresponding cables extended, while the N_{pairs} lowest values have their corresponding cables contracted, where N_{pairs} is a constant predefined parameter that determines how many cable pairs to extend and contract.

All simulations were performed in the MATLAB tensegrity simulation environment [9].

5.5. Robust Training Control

In the first set of simulations, a robust control strategy was trained. A policy to move to three separate goals using the reward function R_2 was selected. A variety of policies with 1, 2, 4, and 8 active cable pairs were trained. Figure 10 shows the resulting fitness curves across 8 statistical trials and their respective standard errors.

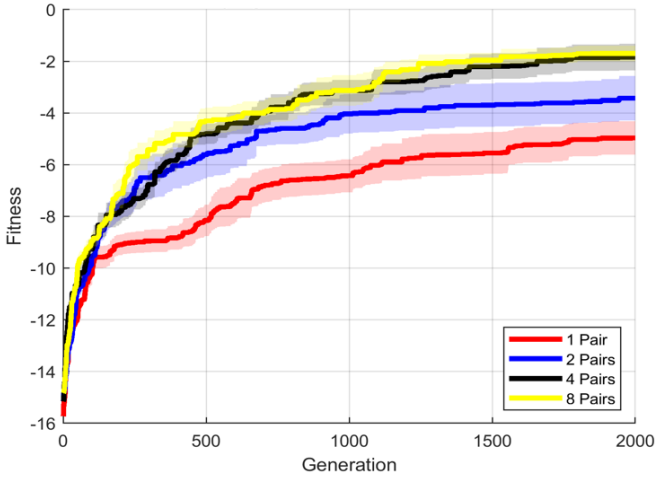

Figure 10: Fitness Policies with Varying Numbers of Active Cables.

There is a minimum number of active cables necessary for a robust controller. The one and two active cable pairs provide some controllability over the tensegrity locomotion, which was improved with four and eight cable pairs.

The next set of simulations determined robustness of the control policies used with cable failure. Each policy tested different variations of the rover with a single cable cut. At the beginning of each episode, a different cable was cut, and the performance of the rover was measured. This led to 78 different performance measurements, one for each cable cut. The team also compared the one active cable pair policies to the eight active cable pair policies and determined the effect of active cable pairs on the robustness to cable failure. Figure 11 shows an average of 8 statistical trials including standard errors.

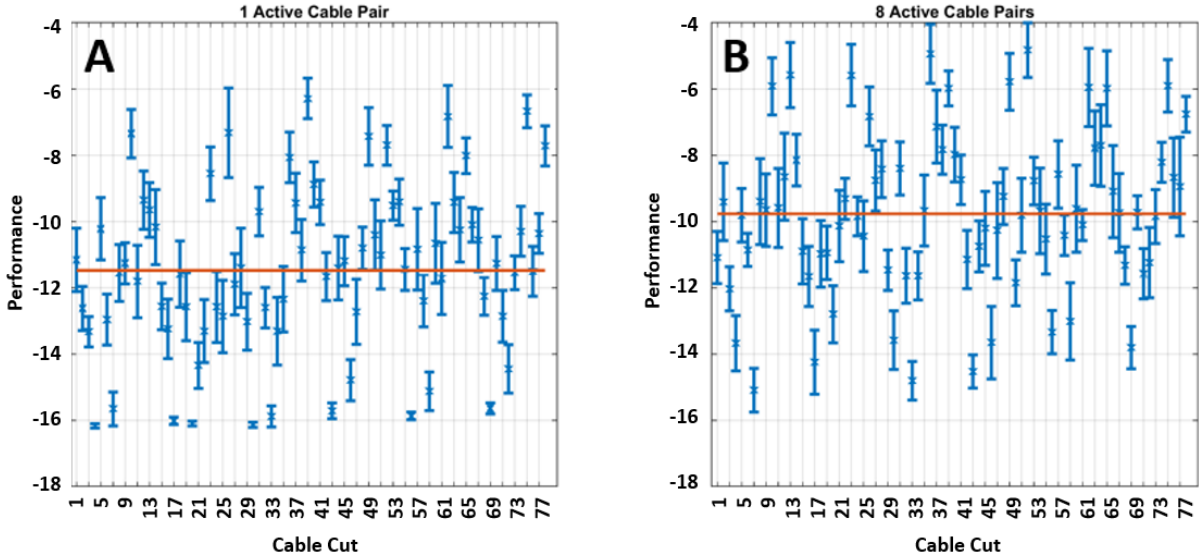

Figure 11: Performance of Controller with Cut Cables.

The x-axes of Figure 11 correspond to the cable that was cut, and the y-axes correspond to performance of the trained controller. The one-active policy performed poorly (a score less than -14) for 13 of the individual cables cut, yet the eight-active policy received a score less than -14 for only four of the cables cut. Two conclusions can be drawn from this. The first is that increasing the number of active cables allows the rover to be more robust to cable failure. The second conclusion is that cables 4 and 7 in the rotational section seen in Figure 4 are crucial to the locomotion of the tensegrity structure. When these cables were cut, the tensegrity structure received a lower performance score.

5.6. Rolling Speed Training

The next set of simulations prioritized training policies to learn quick rolling gaits. To learn this behavior, policies were trained using the R_1 reward with a goal that was too far to reach in the allotted time. Figure 12 shows an example rolling gait. In this figure, the rover starts at the left and moves to the right, with a snapshot roughly every second.

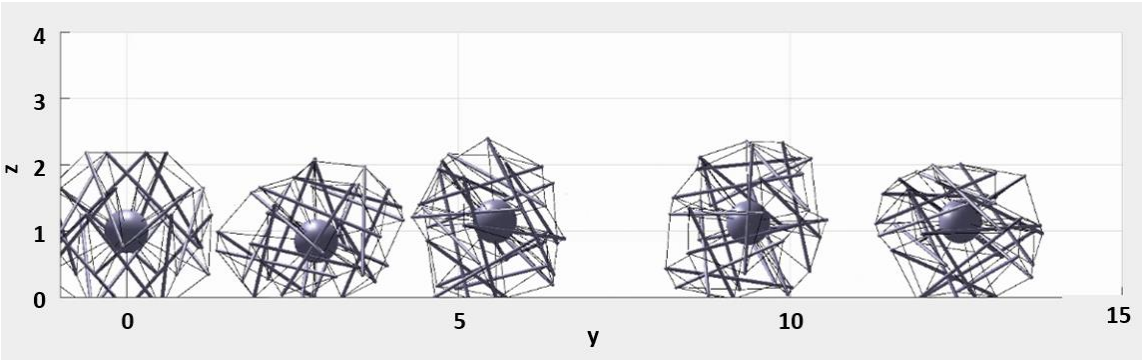

Figure 12: Snapshots of TANDEM Rolling.

The effect of the number of active cable pairs on the speed of the rover was examined. Control policies were trained with 1, 2, 4, 8, and 16 active cables. Each policy was given 20 seconds to reach a goal 30 m from the starting location. The goal was set at a far location so that no rover could reach it, but would have to move quickly to get as close as possible. Figure 13 shows the results from this simulation with 15 statistical trials.

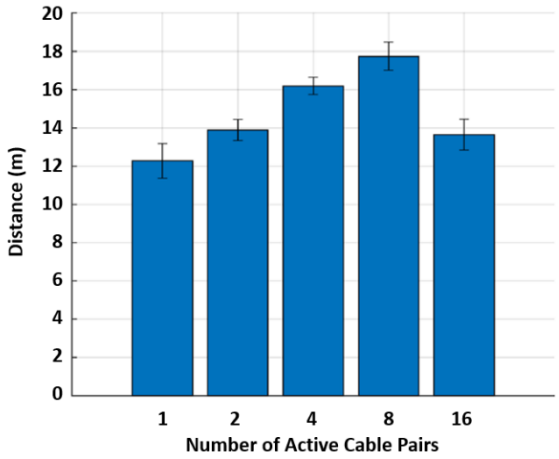

Figure 13: Distance Traveled vs. Number of Active Cable Pairs.

The distance the rover was able to move increased with the number of active cable pairs up to a certain point. At the 16 active cables case, there was a sharp drop off in performance. As the number of active cables increased, the rover showed better control over its shape up until it reached the 16 active cable case, where it became unstable. The resulting policies for 16 cables tend towards the collapse of the rover. The training indicates eight-active cable configuration results in fast and stable gait.

For the last set of simulations, the effects of gravity and payload mass on the speed of the rolling gait learned were tested. The rolling motion learned by the rover consists of a series of leaning and falling forward motions. The speed at which this falling occurs should have a direct impact on the speed of the rover. Simulations were run for the gravities of Titan, the Moon, Mars, and Earth, with payload masses of 10, 50, and 200 kg. Again, the rover was tasked with moving to a goal 30 m away in a span of 20 s. Each case used the same actuator properties. Figure 14 shows the results of 15 statistical runs along with standard error.

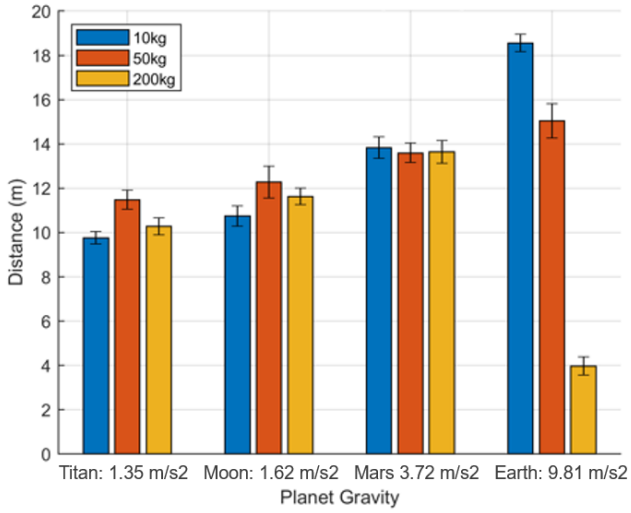

Figure 14: Distance Traveled vs. Gravity and Payload Mass.

The results show a clear correlation between gravity and rolling speed. As gravity rises, the speed of the rover also increases. This makes sense as the rolling motion is a continual falling motion. Interestingly, there was a weak

correlation to mass except for Earth. Here, the mass of the payload has a significant impact on performance. In the case of the 200 kg payload, the gravitational force was large enough to overcome the forces generated by the motors, causing the structure to collapse mid roll. In the 10 kg case, the actuators were able to push the rover faster as they were not pushing against the mass of the payload. For lower gravity environments however, the mass of the payload did not seem to produce a significant effect.

5.7. Potential Control Strategy Improvements

Several paths are present to improve the quality of the controller. Training in a more varied environment with more rugged terrain could help the rover learn robust terrain navigation. Adding to this, the symmetry of the rover could be exploited to improve the quality of the controller. The controller could be reoriented so that it is aligned with the same relative section out of the six sections in the rover.

Another challenge will be the transfer of the policy to a real rover. This is caused by the learner learning a control policy in a simulated environment that does not match the physics of the real world. Training in a higher fidelity training environment such as the NASA Tensegrity Robotics Toolkit could address some of these issues as the physics-based simulator is of higher fidelity than the MATLAB simulator [22]. To improve the transfer to the real rover, noise could be added to the actions in the learning environment [23].

6. Concluding Remarks

The purpose of this research was to expand the application of tensegrity-actuated planetary rovers and demonstrate the TANDEM concept is a viable platform to explore the surface of Titan. Detailed investigation into four major aspects of the mission were reported: entry system analyses, aeroshell design, tensegrity modeling, and locomotion studies. The entry analyses included aerodynamics, aerothermodynamics, trajectory, and mass modeling. The aeroshell was designed using the same ballistic coefficient as the Huygens aeroshell but with the geometry of the Stardust backshell. The structural and TPS masses of the aeroshell were estimated. The master equipment lists for a scaled aeroshell design were reported.

Following the selection of the Stardust aeroshell design, research focused on the specific design of the forward shell. This involved conceptualizing adaptations to accommodate TANDEM and studying optimization techniques for preliminary structural analysis in the near future. It was determined that the estimated impact velocity of the tensegrity rover without a parachute is within the design constraints, eliminating the need for a landing parachute.

Omnidirectional impact studies identified the weak impact orientations for the TANDEM's original topology. Reversing the direction of one of the rigid elements in each section reduced the maximum deceleration experienced by the payload module across all angles of impact. Kinematical, nonlinear optimization form-finding techniques were used to generate several TANDEM configurations. A deployment strategy was developed, including the cable rest-length controller and simulated tension histories for a given shape change.

Lastly, a learning-based control strategy was proposed using neuroevolution to produce efficient and adaptable locomotion gaits. Two different behaviors were trained, one which prioritized a robust controller to move in multiple directions, and the other prioritized speed in a particular direction. Each behavior used a strategy which set a number of cable pairs at a time that actuated together. This way, the tensegrity structure retained its stable, ball-shaped configuration while one cable expanded and the other cable in the pair contracted. The TANDEM concept for the Titan mission is a single, multifunctional tensegrity system which is lightweight, highly packable, and an ideal rideshare candidate on a future mission to Titan.

7. References

[1] Bayandor, J., Schroeder, K., and Samareh, J. (2016). "Lightweight Multifunctional Planetary Probe for Extreme Environment Exploration and Locomotion" NASA Innovative and Advanced Concepts (NIAC) Program Phase I Report, 2016.

[2] Schroeder, K., Samareh, J., and Bayandor, J. (2018). "TANDEM: Tension Adjustable Network for Deploying Entry Membrane," Journal of Spacecraft and Rockets, Vol. 55, No. 6, November-December 2018.

[3] Wercinski, P. F., Smith, B., Yount, B., Kruger, C., Brivkalns, C., Makino, A., Cassell, A., Dutta, S., Ghassemieh, S., and Wu, S. (2017) "ADEPT Sounding Rocket (SR-1) flight experiment overview," 2017 IEEE Aerospace Conference, IEEE, pp. 1–7, 2017.

[4] Samareh, J. (2009). "Multidisciplinary Tool for Systems Analysis of Planetary Entry, Descent, and Landing," NASA/TM-2009-215950, November 2009.

[5] Sepka, S. and Samareh, J. (2015). "Thermal Protection System Mass Estimating Relationships for Blunt-Body, Earth Entry Spacecraft," AIAA Thermophysics Conference, Dallas, TX, June 2015.

[6] Olynick, D., Chen, Y.-K., and Tauber, M. (1999). "Aerothermodynamics of the Stardust Sample Return Capsule," Journal of Spacecraft and Rockets, Vol 36, No. 3, May-June 1999.

[7] Kellas, S. (2016). "Integrated Composite Stiffener Structure (ICoSS) Concept for Planetary Entry Vehicles," AIAA SciTech Conference, San Diego, CA, 2016.

[8] SunSpiral, V., Agogino, A., and Atkinson, D. (2015). "Super Ball Bot – Structures for Planetary Landing and Exploration," Final Phase II Report for the NASA Innovative and Advanced Concepts Program, September 2015.

[9] Friesen, J. (2017). "MATLAB Tensegrity Dynamics Modeling and Simulation Framework," Github, 2017.

[10] Tilbert, A. and Pellegrino, S. (2003). "Review of Form-Finding Methods for Tensegrity Structures," International Journal of Space Structures, Vol. 18, No. 4, pp. 209-223, December 2003.

[11] Vespignani, M., Friesen, J., SunSpiral, V., and Bruce, J. (2018), "Design of SUPERball v2, a Compliant Tensegrity Robot for Absorbing Large Impacts," IEEE/RSJ International Conference of Intelligent Robots and Systems, Madrid, Spain, October 2018.

[12] Gomez, F., Schmidhuber, J., and Miikkulainen, R. (2006). "Efficient Non-linear Control Through Neuroevolution," European Conference on Machine Learning, pp. 654-662, 2006.

[13] Sutton, R. (1997). "On the Significance of Markov Decision Processes," International Conference on Artificial Neural Networks, pp. 273-282, 1997.

[14] Chen, T. and Chen, H. (1995). "Universal Approximation to Nonlinear Operators by Neural Networks with Arbitrary Activation Functions and its Application to Dynamical Systems," IEEE Transactions on Neural Networks, Vol. 6, No. 4, pp. 911-917, 1995.

[15] Such, F, Madhavan, V., Conti, E., Lehman, J., Stanley, K., and Clune, J. (2017). "Deep Neuroevolution: Genetic Algorithms Are a Competitive Alternative for Training Deep Neural Networks for Reinforcement Learning," arXiv:1712.06567, 2017.

[16] Stanley, K. and Miikkulainen, R. (2002). "Evolving Neural Networks Through Augmenting Topologies," Evolutionary Computation, Vol. 10, No. 2, pp. 99-127, 2002.

[17] Gomez. F. (2003). "Robust Non-linear Control Through Neuroevolution," The University of Texas at Austin, 2003.

[18] Iscen, A., Caluwaerts, K., Bruce, J., Agogino, A., SunSpiral, V., and Tumer, K. (2015). "Learning Tensegrity Locomotion Using Open-Loop Control Signals and Coevolutionary Algorithms," Artificial Life, Vol. 21, No. 2, pp. 119-140, May 2015.

[19] Iscen, A., Agogino, A., SunSpiral, V., and Tumer, K. (2013). "Controlling Tensegrity Robots Through Evolution," Proceedings of the 15th annual conference on Genetic and Evolutionary Computation, pp. 1293-1300, July 2013.

[20] Iscen, A., Agogino, A., SunSpiral, V., and Tumer, K. (2014). "Flop and Roll: Learning Robust Goal-Directed Locomotion for a Tensegrity Robot," 2014 IEEE/RSJ International Conference on Intelligent Robots and Systems, Chicago, IL, September 2014.

[21] Miller, B., and Goldberg, D. (1995). "Genetic Algorithms, Tournament Selection, and the Effects of Noise," Complex Systems, Vol. 9, No. 3, pp. 193-212, 1995.

[22] Caluwaerts, K. Despraz, J., Iscen, A., Sabelhaus, A., Bruce, J., Schrauwen, B., and SunSpiral, V. (2014). "Design and Control of Compliant Tensegrity Robots Through Simulation and Hardware Validation," Journal of the Royal Society Interface, Vol. 11, No. 98, September 2014.

[23] Plappert, M., Houthooft, R., Dhariwal, P., Sidor, S., Chen, R., Chen, X., Asfour, T., Abbeel, P., and Andrychowicz, M. (2017). "Parameter Space Noise for Exploration," arXiv:1706.01905, June 2017.